

What is claimed is:

1 1. A method for use in a database system, comprising:
2 receiving a query that specifies an aggregate on distinct values of at least one
3 attribute, the query further specifying grouping on plural grouping sets, the plural
4 grouping sets having at least a first grouping set and a second grouping set;
5 identifying distinct values of the at least one attribute and storing the distinct
6 values of the at least one attribute in a first table;
7 computing aggregates for groups specified by the first grouping set using the first
8 table; and
9 computing aggregates for groups specified by the second grouping set using the
10 first table.

1 2. The method of claim 1, wherein the first grouping set is lower level grouping set
2 than the second grouping set.

1 3. The method of claim 1, wherein identifying the distinct values of the at least one
2 attribute comprises computing a group-by operation on the first grouping set and
3 selecting the attributes of the first grouping set for output.

1 4. The method of claim 3, wherein storing the distinct values of the at least one
2 attribute in the first able comprises storing the distinct values of the at least one attribute
3 in a spool file.

1 5. The method of claim 3, further comprising:
2 using the first table to identify distinct values of the at least one attribute for
3 groups defined by the second grouping set; and
4 storing the distinct values of the at least one attribute for the groups defined by the
5 second grouping set in a second table.

1 6. The method of claim 5, wherein computing aggregates for the groups specified by
2 the second grouping set is based on the second table.

1 7. The method of claim 6, wherein identifying distinct values of the at least one
2 attribute for groups defined by the second grouping set comprises computing a group-by
3 operation on the first table based on the second grouping set and selecting one or more
4 attributes of the second grouping set for output.

1 8. An article comprising at least one storage medium containing instructions that
2 when executed cause a system to:

3 receive a query that specifies an aggregate on distinct values of at least one
4 attribute, the query further specifying grouping on plural grouping sets, the plural
5 grouping sets having at least a first grouping set and a second grouping set;

6 identify distinct values of the at least one attribute and storing the distinct values
7 of the at least one attribute in a first table;

8 compute aggregates for groups specified by the first grouping set using the first
9 table; and

10 compute aggregates for groups specified by the second grouping set using the first
11 table.

1 9. The article of claim 8, wherein the first grouping set is lower level grouping set
2 than the second grouping set.

1 10. The article of claim 8, wherein identifying the distinct values of the at least one
2 attribute comprises computing a group-by operation on the first grouping set and
3 selecting the attributes of the first grouping set for output.

1 11. The article of claim 10, wherein storing the distinct values of the at least one
2 attribute in the first table comprises storing the distinct values of the at least one attribute
3 in a spool file.

1 12. The article of claim 10, wherein the instructions when executed cause the
2 database system to further:

3 use the first table to identify distinct values of the at least one attribute for groups
4 defined by the second grouping set; and

5 store the distinct values of the at least one attribute for the groups defined by the
6 second grouping set in a second table.

1 13. The article of claim 12, wherein computing aggregates for the groups specified by
2 the second grouping set is based on the second table.

1 14. The article of claim 13, wherein identifying distinct values of the at least one
2 attribute for groups defined by the second grouping set comprises computing a group-by
3 operation on the first table based on the second grouping set and selecting one or more
4 attributes of the second grouping set for output.

1 15. A database system comprising:
2 a storage to store a table; and
3 a controller to:

4 receive a query that specifies a calculation of an aggregate on distinct
5 values of an attribute in the table, the query to specify group-by operations on plural
6 grouping sets;

7 in processing the query, compute intermediate values for storage in an
8 intermediate spool; and

9 use the intermediate values in the intermediate spool for computing results
10 of at least two group-by operations on at least two corresponding grouping sets.

1 16. The database system of claim 15, wherein the query comprises a Structured Query
2 Language (SQL) SELECT statement containing a GROUP BY clause specifying multiple
3 grouping sets.

1 17. The database system of claim 15, wherein the query specifies group-by operations
2 on plural grouping sets at multiple grouping levels.

1 18. The database system of claim 15, wherein the controller comprises database
2 management software.

1 19. The database system of claim 18, wherein the database management software
2 comprises plural access modules, and the storage comprises plural storage modules
3 accessible by the plural access modules in parallel.

1 20. The database system of claim 19, further comprising plural processors, the access
2 modules executable on the processors.